

04/09/98

04/09/98

A

+

PTO/SB/05 (2/98) (modified)  
Approved for use through 9/30/2000, OMB 0651-0032  
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

<b>NEW UTILITY PATENT APPLICATION TRANSMITTAL</b>  <i>(only for new nonprovisional applications under 37 CFR 1.53(b))</i>	Attorney Docket Number	3157
	First Named Inventor	Dong-Ying Kuo
	Total Pages in this Submission	33
	Express Mail Label No.	EM207852768US

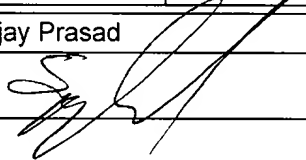
APPLICATION ELEMENTS	ACCOMPANYING APPLICATION PARTS
<p>1. <input checked="" type="checkbox"/> Fee Transmittal Form (in duplicate)  <input checked="" type="checkbox"/> Check Enclosed</p> <p>2. <input checked="" type="checkbox"/> Specification  <i>(preferred arrangement set forth below)</i></p> <ul style="list-style-type: none"> <li><input type="checkbox"/> Descriptive Title of the Invention</li> <li><input type="checkbox"/> Cross Reference(s) to Related Case(s)</li> <li><input type="checkbox"/> Statement Regarding Fed sponsored R &amp; D</li> <li><input type="checkbox"/> Background of the Invention</li> <li><input type="checkbox"/> Brief Summary of the Invention</li> <li><input type="checkbox"/> Brief Description of the Drawing(s)</li> <li><input type="checkbox"/> Detailed Description</li> <li><input type="checkbox"/> Claim or Claims</li> <li><input type="checkbox"/> Abstract of the Disclosure</li> </ul> <p>3. <input checked="" type="checkbox"/> Drawing(s) ( when necessary per 35 USC 113)</p> <p>4. Oath or Declaration</p> <p>a. <input checked="" type="checkbox"/> New Declaration  <input type="checkbox"/> Executed</p> <p>b. <input type="checkbox"/> Copy from a prior application (37 CFR 1.63(d))  <i>(for continuation/divisional with Box 17 completed)</i></p> <p>i. <input type="checkbox"/> DELETION OF INVENTOR(S)  Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b)</p> <p>5. <input type="checkbox"/> Incorporation by Reference (useable if Box 4b is checked). The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.</p>	<p>6. <input type="checkbox"/> Assignment &amp; PTO-1595</p> <p>7. <input type="checkbox"/> Certified Copy of Priority Document(s)  <i>(if foreign priority is claimed)</i></p> <p>8. <input type="checkbox"/> Information Disclosure Statement &amp; PTO-1449  <input type="checkbox"/> Copies of IDS Citation(s)</p> <p>9. <input type="checkbox"/> Preliminary Amendment</p> <p>10. Small Entity Statement  <input type="checkbox"/> New Statement enclosed  <input type="checkbox"/> Statement filed in prior application Status still proper and desired</p> <p>11. <input checked="" type="checkbox"/> Return Postcard</p> <p>12. <input type="checkbox"/> _____</p> <p>13. <input type="checkbox"/> _____</p> <p>14. <input type="checkbox"/> _____</p> <p>15. <input type="checkbox"/> _____</p> <p>16. <input type="checkbox"/> _____</p>

**ADDRESS TO:**  
Assistant Commissioner for Patents  
Box Patent Application  
Washington, D.C. 20231

17. If a **CONTINUING APPLICATION**, check appropriate box and supply the requisite information below and in a preliminary amendment

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No. \_\_\_\_/\_\_\_\_

Prior application information Examiner: \_\_\_\_\_ Group/Art Unit: \_\_\_\_\_

18. CORRESPONDENCE ADDRESS					
NAME	Sanjay Prasad Fenwick & West LLP				
ADDRESS	Two Palo Alto Square				
CITY	Palo Alto	STATE	CA	ZIP CODE	94306
COUNTRY	U.S.A.	TELEPHONE	(650) 858-7291	FAX	(650) 494-1417
Name (Print/Type)	Sanjay Prasad			Registration No. (Attorney/Agent)	36,247
Signature				Date	4/9/98

+

**TOTAL AMOUNT OF PAYMENT**Subtotal (1) + Subtotal (2) + Subtotal (3) = **(\$)790.00**

**Complete if Known**

Application Number	Not Yet Known
Filing Date	April 9, 1998
First Named Inventor	Dong-Ying Kuo
Group Art Unit	Not Yet Known
Examiner Name	Not Yet Known
Attorney Docket Number	3157

**FEE CALCULATION (continued)**

### 3. ADDITIONAL FEES

<u>Large Entity</u> Fee Code/Fee	<u>Small Entity</u> Fee Code/Fee	Fee Description	Fee Due
105/ <b>\$130</b>	205/ <b>\$65</b>	Surcharge - late filing fee or oath	<input type="text"/>
127/ <b>\$50</b>	227/ <b>\$25</b>	Surcharge-late provisional filing fee or cover sheet	<input type="text"/>
147/ <b>\$2,520</b>	147/ <b>\$2,520</b>	For filing a request for reexamination	<input type="text"/>
115/ <b>\$110</b>	215/ <b>\$55</b>	Extension for response within first month <sup>†</sup>	<input type="text"/>
116/ <b>\$400</b>	216/ <b>\$200</b>	Extension for response within second month <sup>†</sup>	<input type="text"/>
117/ <b>\$950</b>	217/ <b>\$475</b>	Extension for response within third month <sup>†</sup>	<input type="text"/>
118/ <b>\$1,510</b>	218/ <b>\$755</b>	Extension for response within fourth month <sup>†</sup>	<input type="text"/>
128/ <b>\$2,060</b>	228/ <b>\$1,030</b>	Extension for response within fifth month <sup>†</sup>	<input type="text"/>
119/ <b>\$310</b>	219/ <b>\$155</b>	Notice of Appeal	<input type="text"/>

**A Duplicate Copy of this authorization is attached**

141/ <b>\$1,320</b>	241/ <b>\$660</b>	Petition to revive unintentionally abandoned application	<input type="text"/>
142/ <b>\$1,320</b>	242/ <b>\$660</b>	Utility Issue Fee (Or Reissue)	<input type="text"/>
143/ <b>\$450</b>	243/ <b>\$225</b>	Design Issue Fee	<input type="text"/>
122/ <b>\$130</b>	122/ <b>\$130</b>	Petitions to the Commissioner	<input type="text"/>
123/ <b>\$50</b>	123/ <b>\$50</b>	Petitions related to provisional applications	<input type="text"/>
126/ <b>\$240</b>	126/ <b>\$240</b>	Submission of Information Disclosure Statement	<input type="text"/>
581/ <b>\$40</b>	581/ <b>\$40</b>	Recording each patent assignment per property (times number of properties)	<input type="text"/>
146/ <b>\$790</b>	246/ <b>\$395</b>	Filing a submission after final rejection (37 CFR 1 129(a))	<input type="text"/>
149/ <b>\$790</b>	249/ <b>\$395</b>	For each additional invention to be examined (37 CFR 1 129(b))	<input type="text"/>

**1. FILING FEE**

**SUBTOTAL (1)** **(\$)790**

<u>Large Entity</u> Fee Code/Fee	<u>Small Entity</u> Fee Code/Fee	Fee Description
103/ <b>\$22</b>	203/ <b>\$11</b>	Claims in excess of 20
102/ <b>\$62</b>	202/ <b>\$41</b>	Independent claims in excess of 3
104/ <b>\$270</b>	204/ <b>\$135</b>	Multiple dependent claim
109/ <b>\$62</b>	209/ <b>\$41</b>	Reissue independent claims over original patent
110/ <b>\$22</b>	210/ <b>\$11</b>	Reissue claims in excess of 20 and over original patent

Other fee (specify)

**SUBTOTAL (3)** (\$).00

(Col. 1)		(Col. 2)		(Col. 3)			
For	No of Existing Claims		Highest No Previously Paid For		Extra**		Fee
TOTAL	20	minus*	20 or	=	0	x	22
INDEP	2	minus*	3 or	=	0	x	82
<input type="checkbox"/> First presentation of multiple dependent claim							

**SUBTOTAL (2)** (\$).00

\*\* If the difference between Col. 1 and Col. 2 is less than zero, then enter "0" in Col. 3

Complete (if applicable)
--------------------------

Date	4/9/98
------	--------

Rev 04/09/98

# Non-Stalled Requesting Texture Cache System and Method

Inventors:

Dong-Ying Kuo

Zhou Hong

Randy Zhao

Roger Niu

Poornachandra Rao

Lin Chen

Jeremy Alves

## BACKGROUND OF THE INVENTION

### Field of the Invention

This invention generally relates to the field of computer graphics systems and more particularly to the generation and processing of textures in computerized graphical images.

### Description of Related Art

To achieve realism in a three-dimensional pixel-based computer-generated graphic image, realistic surface shading must be added to the graphic image in order to simulate real-world surface features such as surface texture. A conventional manner for generating such surface characteristics is texture mapping.

Texture mapping is a relatively efficient technique for creating the appearance of a complex image without the tedium and excessive computational costs of having to directly render three-dimensional details on the surface of a three-dimensional graphic image. In particular, a source (texture) image, which is comprised of individual elements (texels), is

mapped through pixel composition onto the three-dimensional graphic image. For example, to generate a graphic image of an oak table, first the structure of the chair is created. Second, a texture image of oak wood is mapped over the structural surface of the chair. To correctly integrate this oak texture image onto the chair, a predefined number of texels are interpolated to generate each pixel of the surface of the chair.

As is well known in the art, two of the more common texture mapping techniques to effectively interpolate pixels from a texture image are bi-linear filtering and tri-linear filtering. In bi-linear filtering, four texels are used to interpolate each pixel of the graphic image. In tri-linear filtering, eight texels are used to interpolate each pixel of the graphic image.

To avoid potentially overwhelming bandwidth requirements, which are attributable to the use of these texture mapping techniques, mip (*multum in parvo*) mapping often is used to compress the texture image into texture maps of different 'd' levels of resolution. During the scan conversion phase of rasterization, each of these mip maps of the texture map is pre-computed with a specific 'd' level of detail and is separately stored. Depending upon the level of detail of the pixel-based graphic image, a specific 'd' level of the texture map is retrieved for texture filtering. For example, if the graphic image is displayed as a smaller image in the background of a computer-generated graphical scene, the level of detail is low, thereby allowing a lower resolution (e.g., d=4) texture map to be used. With less texels needed for the texture filtering, the lower resolution texture map results in lower system bandwidth requirements.

Alternatively, if the graphic image is displayed as a full-size image in the foreground of a graphic scene, a higher resolution texture map surface (e.g., d=0) is used, thereby raising the overall system bandwidth requirements.

Unfortunately, as discussed in "Hardware Accelerated Rendering of Anti-aliasing Using a Modified A-Buffer Algorithm," Computer Graphics, 307-316, (Siggraph '97 Proceedings), which is incorporated by reference herein, in a real-time environment, where several tens of millions of pixels must be generated every second by these complex texture mapping techniques, high computational demands and tremendous bandwidths for a texture mapping system are unavoidable. For example, with a primary goal of the texture mapping system achieving an average pixel composition rate of one pixel per clock cycle, without the use of caching, one memory block (e.g. four 32-bit texels in a 2 X 2 matrix configuration) must be retrieved every clock cycle. Based upon a clock frequency of approximately 100 MHz, the corresponding bandwidth requirements for achieving such an ideal data transfer rate is approximately 1600 Mbytes per second (MBps).

Due to the ever increasing user expectations of realism in computer generated graphics as well as the demand for economical texture mapping systems, achieving such ideal data transfer rates at an economical price is often elusive. For example, commercially available systems are available that integrate a texture mapping system together with a parallel DRAM memory module sub-system design. Even though such systems offer high data transfer rates, they are economically impractical to the average consumer due to the high price of such designs.

A more inexpensive implementation of the texture mapping system design is integration of the texture mapping system into a graphic accelerator card, which interfaces through an advanced graphics port / peripheral component interconnect (AGP/PCI) interface with a relatively inexpensive general purpose, personal computer. This implementation, however, is confronted with at least two significant performance limitations.

First, due to the AGP/PCI interface only having a bandwidth of approximately 512 MBps, the ability of the texture mapping system to achieve data rates of 1600 MBps is severely curtailed. In particular, such a bandwidth bottleneck results in a relatively significant latency between the system requesting a plurality of texels (memory block) from the memory module and the system receiving back the requested memory block. Second, by issuing these memory block requests in the order that the pixel identifiers are generated by the scan module, the texture mapping system often must page-switch DRAM pages within the memory module to obtain the necessary memory blocks, thereby injecting additional latency into the overall texture mapping system.

To address such performance issues, cost-sensitive texture mapping systems employ a cache module, which is internally located within the system to attempt to minimize the limitations of the AGP/PCI. By retrieving memory blocks prior to needing any of the texels contained within the memory block, the latency attributable to retrieving these texels through the AGP/PCI interface is reduced.

The performance advantages of this local cache design, however, are dependent upon the success ('hit') rate of storing the needed memory blocks within the cache module. For example, when rendering certain shapes (such as long, narrow triangles) as texture primitives, poor memory block locality issues can emerge. Due to cost constraints that require the cache module to be of a limited storage size, the memory block stored within the cache module, which is associated with the beginning of the scan conversion, is no longer locally stored within the cache module by the time the end of a scan conversion span is reached. This situation results in not only an increase in the failure (miss) rate to locate a specific memory block within the cache module, but also in an increase in the need to excessively swap memory blocks into and out of

the cache module. To minimize this locality problem, texture mapping systems require accurate synchronization between the memory module and the cache module without dramatically increasing the cost of the system.

What is needed is a texture mapping system and method, which can effectively address the issues of transmitting as many texel requests as early as possible, synchronizing the processing of texel requests and the corresponding texels and avoiding bandwidth limitations in a cost-effective manner.

## SUMMARY OF THE INVENTION

In accordance with the system and method of the present invention, an accelerated caching system is invented to increase the operation concurrency between a cache module and a memory module based upon efficient synchronization of a first in, first out scheme. The caching system may comprise a scan module, a memory module, a texture cache module, a texture filter module and a pixel processing unit.

The scan module generates pixel identifiers corresponding to both specific pixels in a three-dimensional graphic image and memory blocks of texels within the memory module. The texture cache module strips each of these pixel identifiers into its corresponding memory block (MB) identifier and texel identifier subcomponents. Based upon a first in, first out scheme, the texture cache module utilizes each MB identifier to determine whether a memory block corresponding to the MB identifier has already been locally stored within the cache module. If the memory block is already locally cached, the texture cache module synchronizes the transmission of the texels within the memory block to the texture filter module. If the memory block is not already locally cached within the texture cache module, the texture cache module

retrieves the needed memory block from the memory module and stores the memory block in a specific location within the texture cache module. When these texels are needed for pixel composition, the texture cache module transmits the texels to the texture filter module.

This system and method ensures that a high performance and cost effective first in, first out caching scheme is achieved by increasing the operation concurrency between the cache module and the memory module through efficient synchronization.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

Figure 1 is a high level illustration of a preferred embodiment of the texture mapping system of the preferred embodiment of the present invention.

Figure 2 is a detailed illustration of the texture cache module of the preferred embodiment.

Figure 3 is a detailed illustration of the address module of the preferred embodiment.

Figure 4 is an illustration of the controller module of the preferred embodiment.

Figure 5 is an illustration of the translation of the memory block storage scheme to the cache block storage scheme in the preferred embodiment.

### **DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS**

A preferred embodiment of the present invention is now described with reference to figures where like reference numbers indicate identical or functionally similar elements and the left most digit of each reference number corresponds to the figure in which the reference number is first used.



Figure 1 is an overview of the system of the preferred embodiment of the present invention for improving the performance of caching texels for pixel composition. The system 100 includes a memory module 140, a scan module 160, a texture cache module (TCM) 130, a texture filter module (TFM) 180 and a pixel processing unit (PPU) 190. In this embodiment, the memory module 140 is comprised of dynamic random access memory (DRAM), which stores each 4 X 4 matrix of texels in one of a plurality of memory blocks.

The scan module 160 coupled to the TCM 130 transmits a plurality of pixel identifiers 510 (shown in Figure 5) to the TCM 130. The TCM 130 utilizes each of these pixel identifiers 510 to ensure that a memory block corresponding to the pixel identifier 510 has been retrieved from the memory module 140 and locally cached within the TCM 130. Once the memory blocks are locally cached, each texel within the memory block corresponding to the pixel identifier 510 is transmitted to the TFM 180. The TFM 180 filters the received texels into resultant texels, which then are transmitted to the PPU 190 for pixel composition

Figure 2 is a more detailed illustration of the TCM 130 of the preferred embodiment of the present invention. The TCM 130 comprises a cache module 220, a data formatter 255, a retrieval buffer 260, a request buffer 250, an address module 210, a buffer module 240, and a controller module 290.

The cache module 220, which in this embodiment is comprised of SRAM having a 4 way - 4 set association, locally stores a subset of the memory blocks stored within the memory module 140. As is illustrated in figure 5, such an association results in the cache module 220 becoming organizationally divided into four ways (e.g., ways 0-3) with each way having four quadrants (sets). Each set having a unique cache block address (e.g., a number between 0 and 15) represents a 4 x 4 matrix of texels (cache block).

The address module 210, which is coupled to the scan module 160 and the cache module 220, receives from the scan module 160 a pixel identifier 510 ( $u.u_f, v.v_f, d.d_f$ ), which includes an integer portion ( $u._, v._, d._$ ) and a fractional portion ( $_.u_f, _.v_f, _.d_f$ ), for each scanned pixel. From the integer portion of each pixel identifier 510, the address module 210 generates a memory block (MB) identifier and a texel identifier 520. The address module 210 utilizes the MB identifier 530 to identify a memory block within the memory module 140. The address module 210 utilizes the texel identifier 520 to locate a specific texel within that memory block. The address module 210 relies upon the fractional portion of the pixel identifier 510 as an interpolation factor for linearly interpolating the texels corresponding to the texel identifier 520 for pixel composition.

Figure 5, in conjunction with figure 2, better illustrates the use of each of these identifiers in the preferred embodiment. For bi-linear filtering, the scan module 160 generates four, 11 bit [10:0] pixel identifiers 510 for every pixel of the three-dimensional graphic image. The address module 210 strips each integer portion of the pixel identifier 510 into a 9 bit MB identifier 530 and a 2 bit texel identifier 520. The MB identifier 530, which represents a 4 X 4 texel memory block within the memory module 140, is the highest 8 bits [10:2] of the integer portion. The texel identifier 520, which represents an orthogonal coordinate for a texel within this memory block, is the lower 2 bits [1:0] of the integer portion.

The address module 210 utilizes the MB identifier 530 to determine whether the memory block corresponding to the MB identifier 530 already has been locally stored within a cache block of the cache module 220. If the address module 210 determines that the memory block is already locally cached at a specific cache block address (hit), the address module 210 generates and transmits a transmit tag signal to the buffer module 240. The transmit tag signal includes the

cache block address, the texel identifier 520, the interpolation factors and a hit marker. The TCM 130 relies upon the hit marker to identify transmit tag signals corresponding to memory blocks, which are already locally cached within the cache module 220.

If the memory block is not already locally cached (miss), the address module 210 initiates a request for retrieval of the memory block from the memory module 140 by transmitting a retrieval signal to the request buffer 250, which operates as a FIFO queue. In this embodiment, the request buffer 250 can receive up to 4 retrieval signals at the same time. The retrieval signal includes the MB identifier 530 and a linear memory module address. Since the retrieval signal cannot be transmitted to the memory module 140 any faster than the AGP/PCI interface (not shown) will allow, the request buffer 250 continuously transmits the retrieval signals to the memory module 140 in the order the corresponding pixel identifiers 510 were received by the address module 210. The AGP/PCI interface, however, limits the transmission rate of the request buffer 250 to an approximate maximum rate of 512 MBps.

The memory module 140 responds to the receipt of the retrieval signals by transmitting the memory block, which corresponds to the MB identifier 530, back to the TCM 130. However, due to cost constraints, which limit the storage capacity of the cache module 220 to a finite size, not all memory blocks retrieved from the memory module 140 can be immediately stored within the cache module 220. Instead, until the cache block associated with the memory block is available to store the memory block, the retrieval buffer 260, which also operates as a FIFO queue, temporarily stores the received memory blocks in the order received. In this embodiment, the retrieval buffer 260 receives the 256-bit memory block within 2 clock cycles of a system clock.

In addition to transmitting the retrieval signal to the request buffer 250, the address module 210 also generates and transmits a retrieval tag signal and a transmit tag signal to the buffer module 240. The retrieval tag signal includes a tag reference counter value, which will be discussed with regard to figure 3, and a cache block address, which designates a specific cache block within the cache module 220 where the retrieved memory block corresponding to the MB identifier 530 will be stored. The address module 210 determines the specific cache block address for each retrieval tag signal by relying upon a least recently used (LRU) algorithm. The transmit tag signal includes the cache block address, the texel identifier, interpolation factors and a miss marker.

The controller module 290 relies upon the retrieval tag signals in the buffer module 240 to control the transmission of the memory blocks from the retrieval buffer 260 to the cache block associated with the memory block. As will be discussed in more detail with regard to figure 4, when the controller module 290 determines that the cache block is available for the associated memory block at the top of the retrieval buffer 260, the controller module 290 triggers the transmission of the memory block to the data formatter 255. The data formatter 255 transforms this memory block from its original format into a cache module format, which in this embodiment is a conventional ARGB format. The data formatter 255 then transmits each of these reformatted memory blocks to the cache module 220 where the memory blocks are stored in their associated cache blocks.

The controller module 290 relies upon the interpolation information isolated by the address module 210 to determine when to trigger the transmission to the TFM 180 of each texel associated with each transmit tag signal from the cache module 220. Since a filtering technique, such as bi-linear filtering, requires more than one texel to perform pixel composition and there is

the possibility that each of these texels may be located in a different cache block due to the 4 way - 4 set organization of the cache module 220, up to four transmit tag signals must be transmitted to the cache module 220 for every composition of a pixel.

Figure 3 is a more detailed illustration of the address module 210 and the buffer module 240. The address module 210 includes an address controller 310, a tag module 320, and a tag reference counter (TRC) 330. The buffer module 240 comprises a cache buffer 335 and a transmit buffer 345. Each buffer 335 and 345 can receive up to 4 tag signals at the same time. The tag module, which is coupled to the address controller 310, stores associations between each cache block address and the MB identifier 530 corresponding to the memory block stored within that cache block. In this embodiment, the tag module 320 is a table, which is edited by the address controller 310.

The address controller 310 is coupled to the request buffer 250, the tag module 320, the TRC 330, the cache buffer 335, the transmit buffer 345 and the scan module 160. In this preferred embodiment, the address controller 310 is a state machine, which either is a software program executed by a central processing unit (CPU) (not shown) of an Intel-based personal computer, or alternatively is comprised of hardware organized to implement a state machine. Upon receipt of the pixel identifiers 510 from the scan module 160, the address controller 310 compares the MB identifier 530 portion of each pixel identifier 510 with the MB identifiers 530 stored within the tag module 320. The TRC 330 records the consecutive number of times the address controller 310 achieves a match (hit) between these two MB identifiers 530. In this embodiment, the TRC 330 is an 8-bit counter, which initially is set to 0.

If the address controller 310 matches a received MB identifier 530 with an MB identifier 530 stored within the tag module 320, the address controller 310 generates the before-mentioned

transmit tag signal containing the hit marker, the texel identifier 520, corresponding to the MB identifier 530, and the cache block address, associated with the MB identifier 530. The address controller 310 also increments the TRC 330 by a value of 1.

If the address controller 310 is unable to match a received MB identifier 530 with one of the MB identifiers 530 stored within the tag module 320, the address controller 310 generates a retrieval tag signal, which includes the MB identifier 530 and a linear memory module address. In addition, a transmit tag signal and a retrieval tag signal are generated and transmitted to the buffer module 240. The retrieval tag signal includes the current TRC value and a cache block address, which, as previously discussed, is associated by the address controller 310 with the unmatched MB identifier 530 through the use of the LRU algorithm. This cache block address later is used as the location within the cache module 220 where the retrieved memory block corresponding to the unmatched MB identifier 530 will be stored. Unlike the 'hit' situation, in the 'miss' situation the TRC 330 is not incremented by 1. Rather, the current value of the TRC 330 is incorporated into the retrieval tag signal and the TRC 330 is set back to 1. This TRC value represents the number of consecutive 'hits' before the address controller 310 identified the 'miss' situation. The transmit tag signal for the 'miss' situation includes the texel identifier 520, the cache block address corresponding to the texel identifier 520, and a 'miss' marker. The 'miss' marker is used to identify those transmit tag signals, which correspond to memory blocks that must be retrieved from the memory module 140. Both the retrieval tag signal and the transmit tag signal then are appended to the end of the cache buffer 335 and the transmit buffer 345, respectively.

Figure 4 illustrates a detailed view of the controller module 290 of the TCM 130. The controller module 290 includes a retrieval controller 410, a transmit controller 430, a data

reference counter (DRC) 440 and a status module 450. The DRC 440 records the number of consecutive transmit tag signals having a 'hit' marker, which are transmitted to the cache module 220. In this embodiment, the DRC 440 is an 8-bit counter and the status module 450 is a table where each cache block address is associated with a logic state (e.g., "read only" logic state or "write only" logic state).

The transmit controller 430, which is coupled to the transmit buffer 350, the DRC 440 and the status module 450, controls when the texels within a memory block, which correspond to the transmit tag signals at the top of the transmit buffer 350, are further processed by the TCM 130. In the preferred embodiment, the transmit controller 430 is a state machine, which is either a software program executed by a CPU of an Intel-based personal computer, or alternatively comprises hardware components organized to implement such a state machine. If the transmit tag signals relating to a specific pixel both contain a 'hit' marker and are the next signals to be transmitted by the transmit buffer 350 as well as the logic state in the status module 450 of the corresponding cache blocks is set to the 'read only' state, the transmit controller 430 triggers the transmit buffer 350 to transmit the transmit tag signals to the cache module 220. The cache module 220 then transmits the texels corresponding to the transmit tag signals to the TFM 180. At approximately the same time, the transmit controller 430 also increments by one the DRC 440, which is initially set to a value of 0. If at least one of the transmit tag signals, which is at the top of the transmit buffer 350 and is associated with a specific pixel identifier 510, contains a 'miss' marker, the transmit controller 430 does not immediately trigger the transmit buffer 350 to transmit the transmit tag signals to the cache module 220. Rather, the retrieval controller 410 will change the logic states of the cache blocks from the 'read only' to the 'write only' state and will trigger the retrieval of the memory block from the retrieval buffer 260. Upon detecting the

change of the logic state to 'write only', the transmit controller 430 changes the logic state back to the 'read only' state, increments the DRC by 1 and triggers the transmit buffer 350 to transmit the transmit tag signals, which are needed for pixel composition, to the cache module 220.

The retrieval controller 410 is coupled to the retrieval buffer 260, the cache buffer 340, the DRC 440 and the status module 450. In the preferred embodiment, the retrieval controller 410 is a state machine, which is a software program executed by a CPU of an Intel-based personal computer, or alternatively comprises hardware organized to implement such a state machine. The retrieval controller 410 is responsible for synchronizing the writing of memory blocks within the cache block 220 by changing the logic state associated with a specific cache block within the status module 450 from 'read only' to 'write only'. In particular, the retrieval controller 410 waits for the memory block corresponding to the retrieval tag signal to be received by the retrieval buffer 260 from the memory module 140. Once received, the retrieval controller 410 compares the current DRC value with the TRC value within the retrieval tag signal. When this TRC value is greater than 0 and less than the DRC value, the retrieval controller 410 pauses. When the DRC value is greater than or equal to the TRC value and the retrieval buffer 260 is not empty, the retrieval controller 410 sets the logic state corresponding to the cache block address to 'write only' and sets the DRC value back to 0.

As discussed above, the changing of the logic state to 'write only' will result in the retrieval controller 410 triggering the cache buffer 340 to transmit the retrieved memory block at the top of the retrieval buffer 260 to the data formatter 255 for reformatting. The data formatter 255 then transmits the reformatted memory block to the cache module 220 where the memory block is stored in the cache block corresponding to the retrieval tag signal. Once the memory block is stored within the cache module 220, the transmit controller 430 responds to the change



in the logic state to 'write only' by changing the logic state of the cache block back to 'read only'. Even though this entire technique may have the risk of possibly overwriting memory blocks containing texels, which may be needed at a later time, the amount of latency avoided by this technique far outweighs the slight delay that is attributable to having to re-retrieve these  
5 memory blocks.

The foregoing description of the preferred embodiment has been presented for purposes of illustration and description. It is not intended to be exhaustive nor to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. For example, tri-linear filtering, rather than bi-linear filtering can be supported by an  
10 alternative embodiment of the present invention. The preferred embodiment was chosen and described to best explain the principles of the invention and its practical application to thereby enable others skilled in the art to best utilize the invention in various embodiments and with various modifications as are suited to the particular use contemplated. The system and method, however, can be used in any other situation where a bandwidth bottleneck hinders the efficient  
15 processing of large amounts of data. It is intended that the scope of the invention be defined by the claims and their equivalents.

## CLAIMS

What is claimed is:

1. A non-stalled requesting system comprising:

a first storage device for storing a plurality of texels;

a second storage device coupled to the first storage device for storing a first texel of the plurality of texels;

a third storage device for storing an association between a first identifier signal and the first texel;

a first logic device coupled to the second storage device and the third storage device for responding to a second identifier signal matching the first identifier signal by triggering the transmission of the first texel from the second storage device and for responding to the second signal not matching the first identifier signal by retrieving a second texel of the plurality of texels from the first storage device; the second texel having an association with the second identifier signal.

2. The system of claim 1 further comprising:

a second logic device coupled to the first logic device for synchronizing the transmission of the first texel from the second storage device.

3. The system of claim 1 further comprising:

a third logic device coupled to the first logic device for synchronizing the storing of the second texel in the second storage device.

1 4. The system of claim 2 further comprising:

2 a fourth storage device coupled to second logic device for storing an association between  
3 the second storage device and a plurality of logic states.

1 5. The system of claim 1 further comprising:

2 a first counter coupled to the first logic device for responding to the second identifier  
3 signal matching the first identifier signal by changing a first value to a second value.

1 6. The system of claim 2 further comprising:

2 a second counter coupled to the second logic device for responding to the transmission of  
3 the first texel by the second storage device by changing a third value to a fourth value.

1 7. The system of claim 5 further comprising:

2 a second counter for responding to the transmission of the first texel by the second  
3 storage device by changing a third value to a fourth value; and  
4 a third logic device coupled to the first logic device and to the second counter for  
5 responding to the fourth value having a value at least as large as the second value by triggering  
6 the storing of the second texel in the second storage device.

1 8. The system of claim 1 further comprising:

2 a second logic device coupled to the second storage device for controlling the  
3 transmission of the first texel from the second storage device;  
4 a third logic device coupled to the second storage device for triggering the storing of the  
5 second texel in the second storage device; and

6 a third storage device coupled to the second logic device and the third logic device for  
7 synchronizing the transmission of the first texel and the storing of the second texel.

1 9. The system of claim 3 further comprising:

2 a first buffer coupled to the first storage device, the second storage device and the third  
3 logic device for delaying the storing of the second texel in the second storage module until the  
4 third logic device triggers the storing of the second texel in the second storage device.

1 10. The system of claim 9 wherein the first buffer is a first in, first out queue.

1 11. A method for requesting texels, the method comprising the steps of:

2 storing the plurality of texels in a first storage device;  
3 storing a first texel of the plurality of texels in a second storage device;  
4 storing an association between a first identifier signal and the first texel;  
5 responding to a second identifier signal matching the first identifier signal by triggering  
6 the transmission of the first texel from the second storage device; and  
7 responding to the second identifier signal not matching the first identifier signal by  
8 retrieving a second texel of the plurality of texels, having an association with the second  
9 identifier signal, from the first storage device.

1 12. The method of claim 1 comprising the additional step of:

2 synchronizing the transmission of the first texel from the second storage device.

1 13. The method of claim 1 comprising the additional step of:

2 synchronizing the storing of the second texel in the second storage device.

1 14. The method of claim 12 comprising the additional step of:  
2 storing an association between the second storage device and a plurality of logic states.

1 15. The method of claim 11 comprising the additional step of:  
2 responding to the second identifier signal matching the first identifier signal by changing  
3 a first value to a second value.

1 16. The method of claim 12 comprising the additional step of:  
2 responding to the transmission of the first texel by the second storage device by changing  
3 a third value to a fourth value.

1 17. The method of claim 15 comprising the additional step of:  
2 responding to the transmission of the first texel by the second storage device by changing  
3 a third value to a fourth value; and  
4 responding to the fourth value having a value at least as large as the second value by  
5 triggering the storing of the second texel in the second storage device.

1 18. The method of claim 11 comprising the additional step of:  
2 controlling the transmission of the first texel from the second storage device;  
3 triggering the storing of the second texel in the second storage device; and  
4 synchronizing the transmission of the first texel and the storing of the second texel.

1 19. The method of claim 13 comprising the additional step of:  
2 delaying the storing of the second texel in the second storage module until the third logic  
3 device triggers the storing of the second texel in the second storage device.

1 20. The method of claim 19 wherein the delay the storing of the second texel comprises the  
2 additional step of:  
3 transmitting the second texel, which was received prior to a third texel, before  
4 transmitting the third texel.

## ABSTRACT

A caching system for increasing the operation concurrency between a cache module and a memory module by comparing received memory block identifiers, which correspond to texels  
5 needed for pixel composition, with memory block identifiers corresponding to texels locally stored within the cache module. If the received memory block identifiers match the memory block identifiers corresponding to locally cached texels, the system transmits these texels to a texture filter unit for pixel composition. If the received memory block identifiers do not match  
10 memory block identifiers corresponding to the locally cached texels, the system retrieves these texels from the memory module as fast as possible and then updates the cache module with the new texels. A plurality of first in, first out buffers are used to assist a controller module with synchronizing the transmission of the texels from the cache module and the overwriting of the texels received from the memory module into the cache module. Once locally cached, these texels also are transmitted to the texture filter unit for pixel composition.

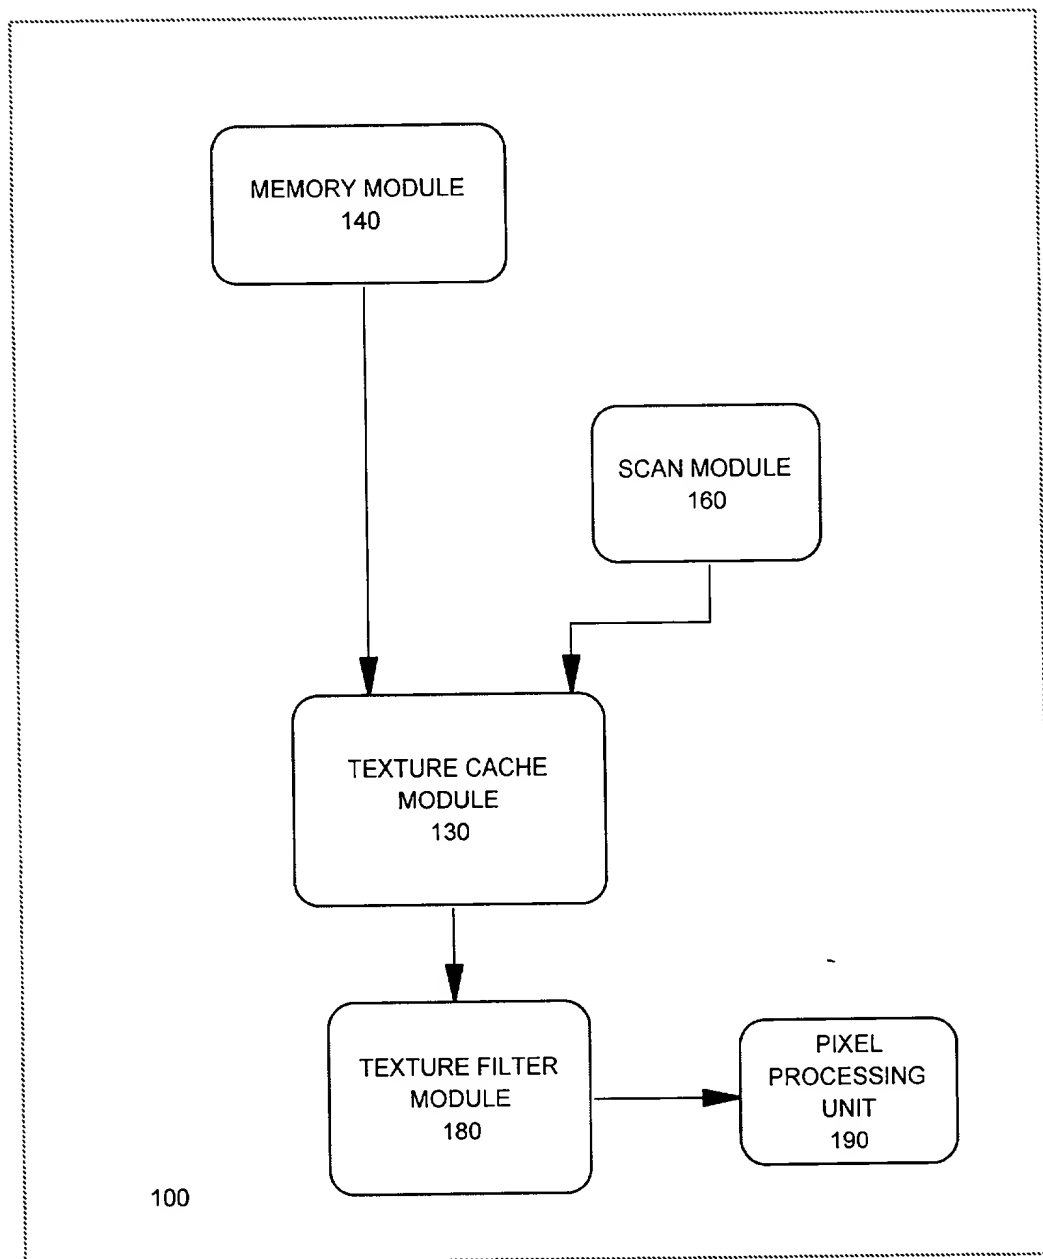


FIGURE 1



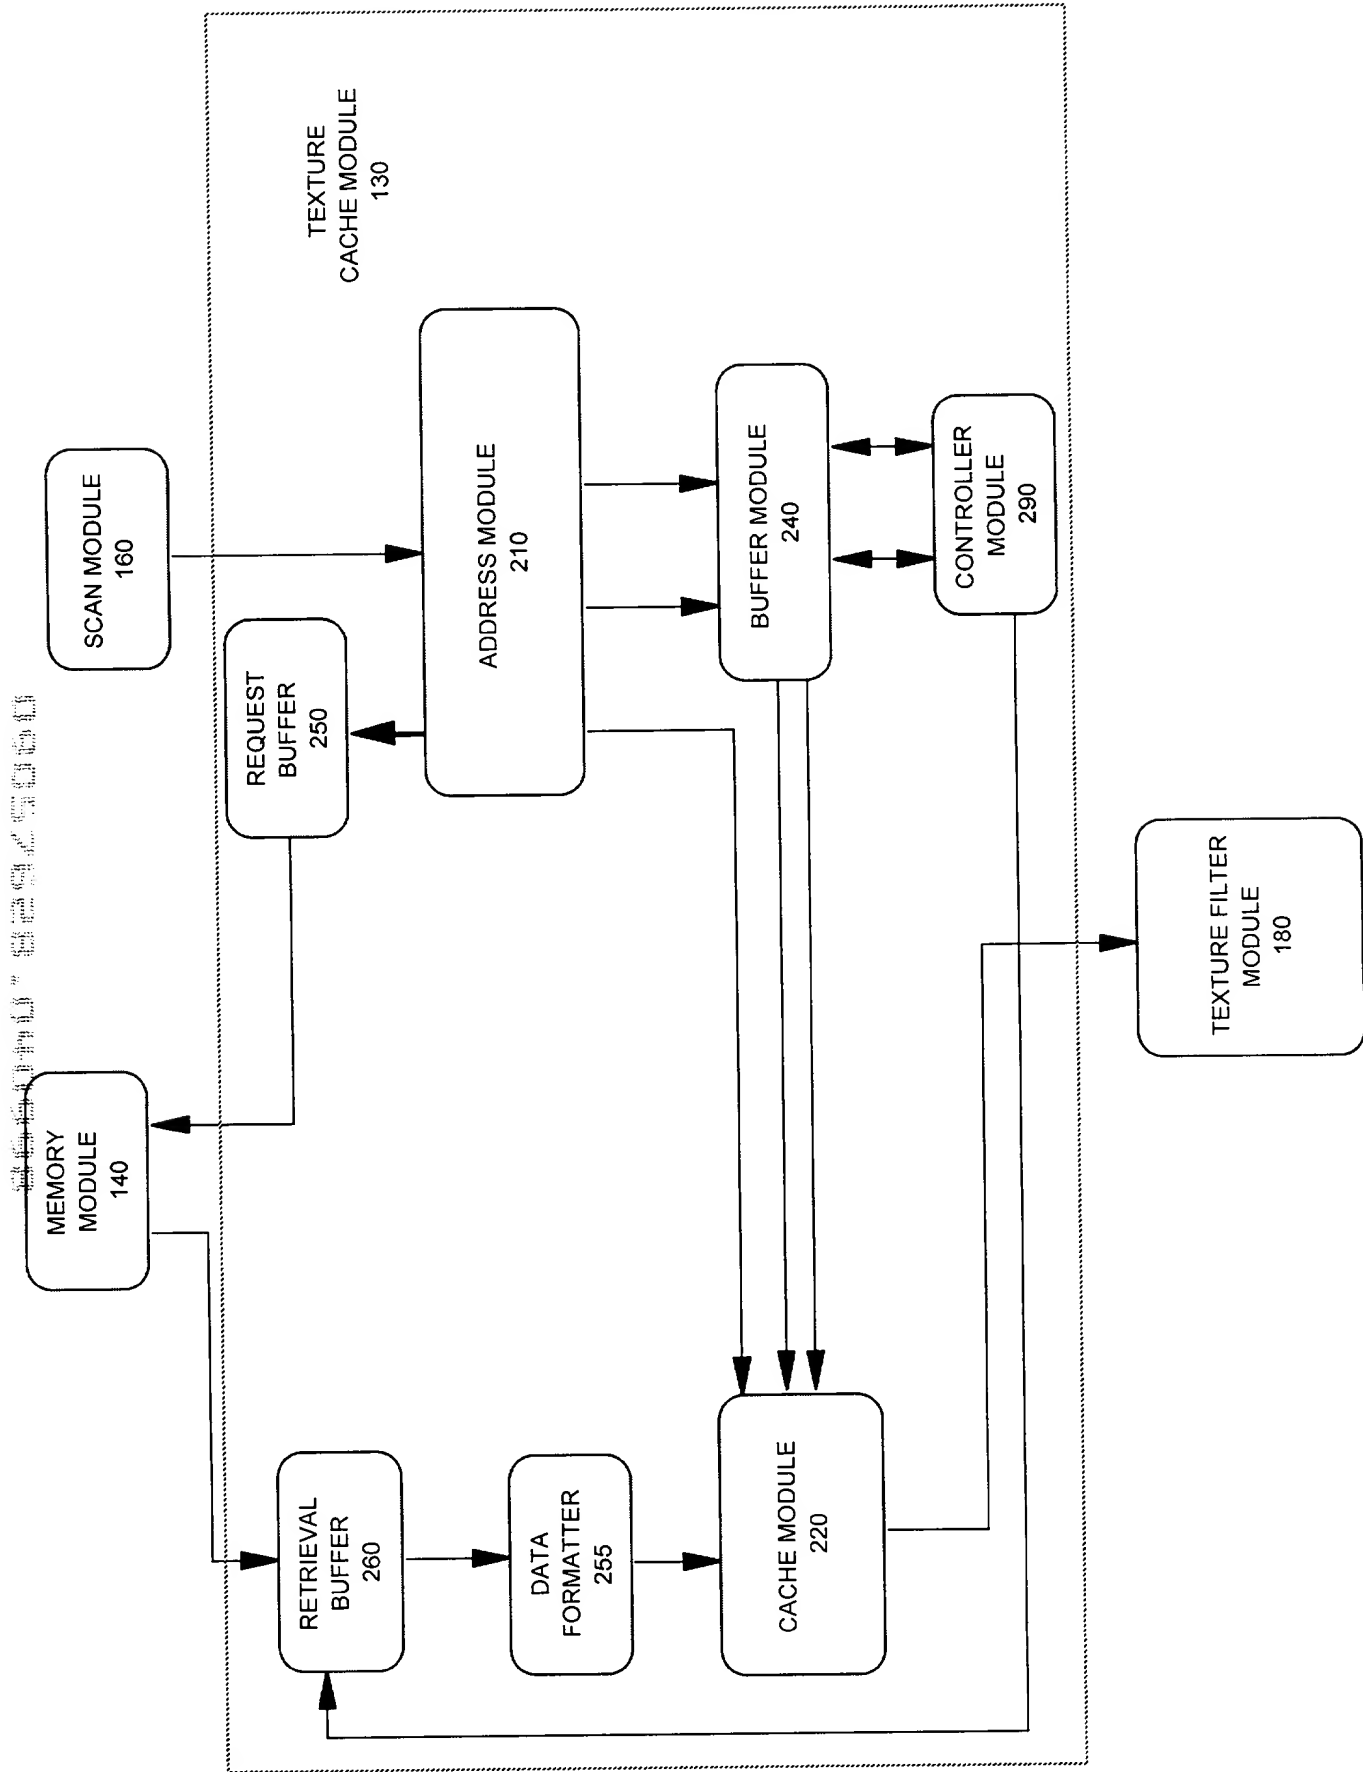


FIGURE 2

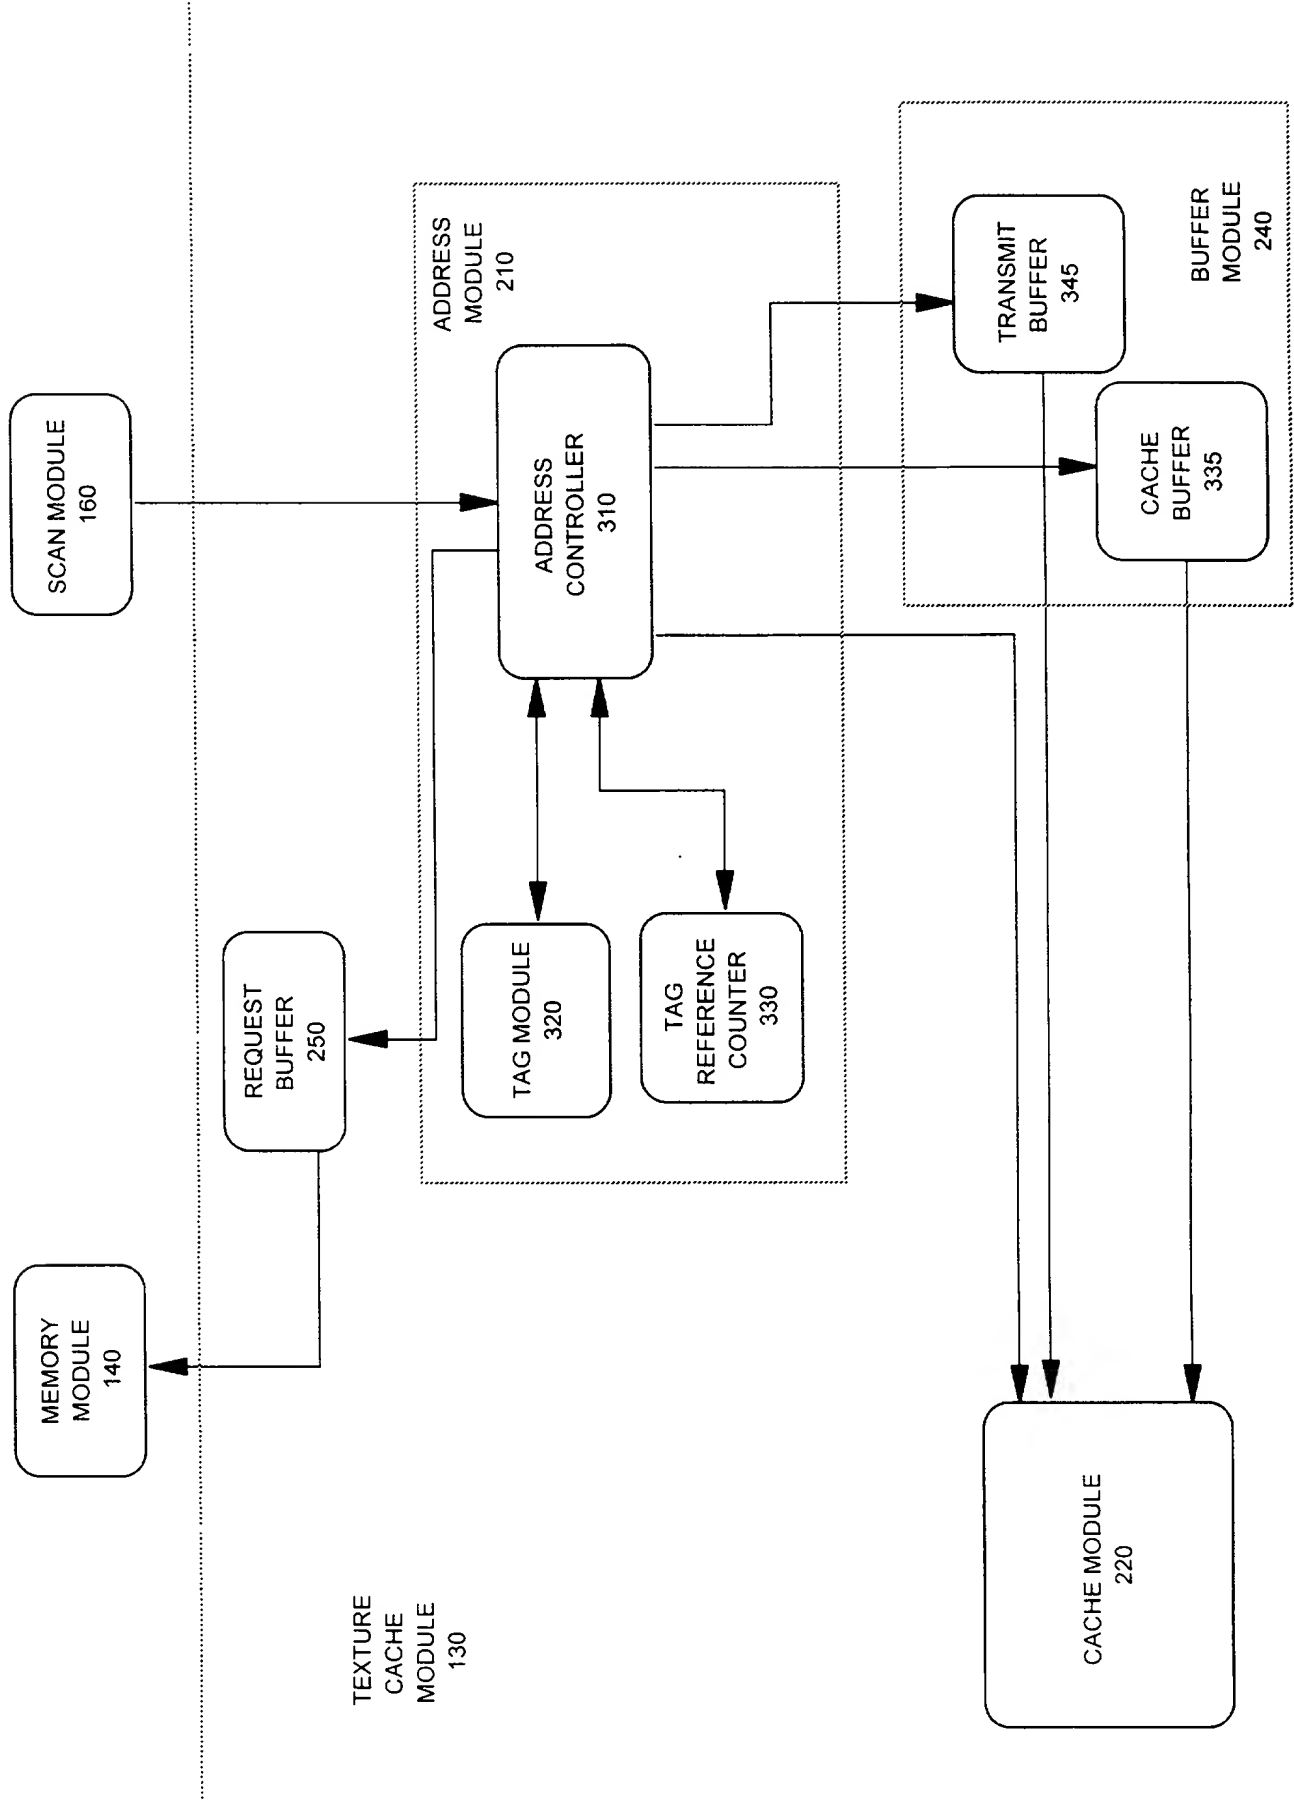


FIGURE 3

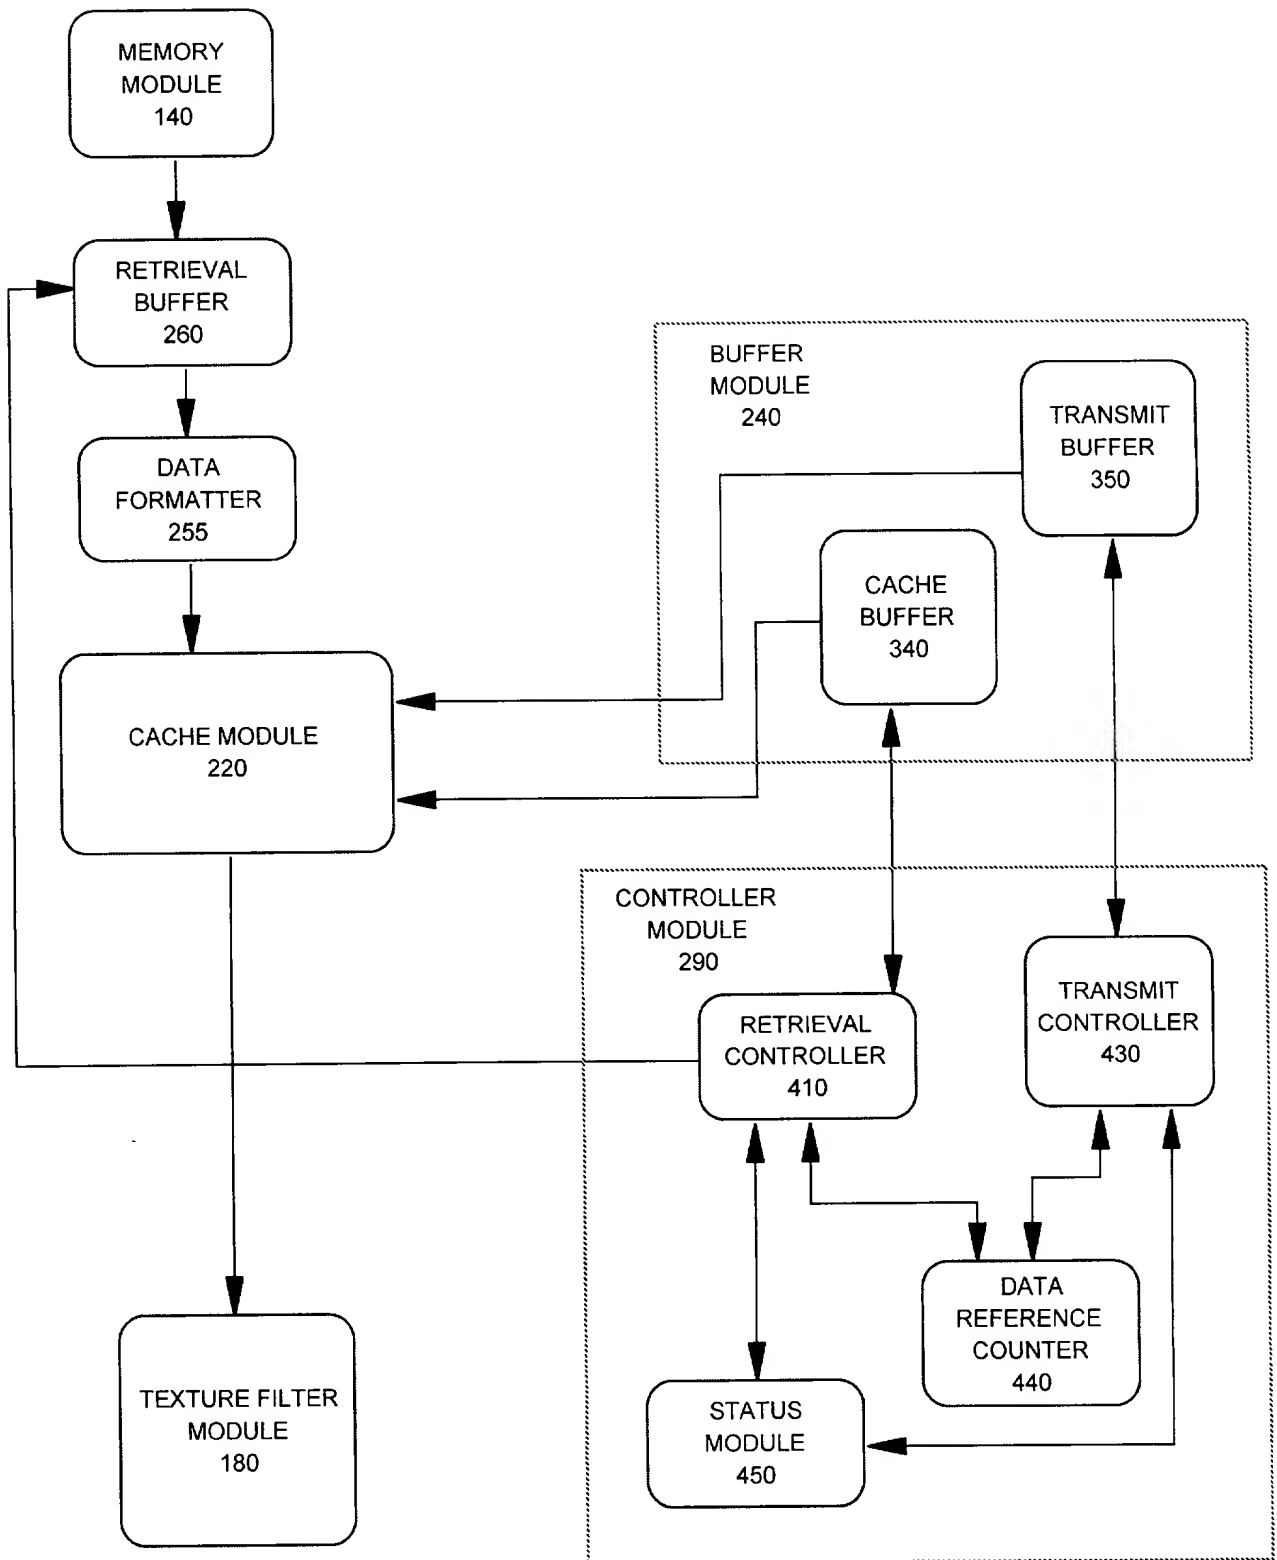


FIGURE 4




FIGURE 5

<b>0010/PTO</b> Rev. 6/95  <b>U.S. Department of Commerce</b> Patent and Trademark Office  <b>DECLARATION FOR UTILITY OR DESIGN PATENT APPLICATION</b>	Attorney Docket Number	3157
	First Named Inventor	Dong-Ying Kuo
	COMPLETE IF KNOWN	
	Application Number	Not Yet Known
	Filing Date	April 9, 1998
	Group Art Unit	Not Yet Known
	Examiner Name	Not Yet Known
<input checked="" type="checkbox"/> Declaration Submitted with Initial Filing      OR <input type="checkbox"/> Declaration Submitted after Initial Filing		

As a below named inventor, I hereby declare that:

My residence, post office address, and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

**NON-STALLED REQUESTING TEXTURE CACHE SYSTEM AND METHOD**

the specification of which

(Title of the Invention)

☒ is attached hereto

OR

☐ was filed on (MM/DD/YYYY) [ ] as United States Application Number or PCT International Application Number [ ] and was amended on (MM/DD/YYYY) [ ] (if applicable).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment specifically referred to above.

I acknowledge the duty to disclose information which is material to patentability as defined in Title 37 Code of Federal Regulations, § 1.56

I hereby claim foreign priority benefits under Title 35, United States Code § 119 (a)-(d) or § 385(b) of any foreign application(s) for patent or inventor's certificate, or § 365 (a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or of any PCT international application having a filing date before that of the application on which priority is claimed

Prior Foreign Application Number(s)	Country	Foreign Filing Date (MM/DD/YYYY)	Priority Not Claimed	Certified Copy Attached?	
				YES	NO
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
			<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

☐ Additional foreign application numbers are listed on a supplemental priority sheet attached hereto:

I hereby claim the benefit under Title 35, United States Code § 119(e) of any United States provisional application(s) listed below.

Application Number(s)	Filing Date (MM/DD/YYYY)	<input type="checkbox"/> Additional provisional application numbers are listed on a supplemental sheet attached hereto

<b>DECLARATION</b>	<b>Page 2</b>
--------------------	---------------

I hereby claim the benefit under Title 35, United States Code § 120 of any United States application(s), or § 365(c) of any PCT international application designating the United States of America, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT international application in the manner provided by the first paragraph of Title 35, United States Code § 112. I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations § 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application.

U.S. Parent Application Number	PCT Parent Number	Parent Filing Date (MM/DD/YYYY)	Parent Patent Number (if applicable)

☐ Additional U.S. or PCT international application numbers are listed on a supplemental priority sheet attached hereto.

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith:

Name	Registration Number	Name	Registration Number
<b>Greg T. Sueoka</b>	<b>33,800</b>		
<b>Sanjay Prasad</b>	<b>36,247</b>		
<b>Leif Sloan</b>	<b>37,942</b>		

☐ Additional attorney(s) and/or agent(s) named on a supplemental sheet attached hereto.

Please direct all correspondence to:

**Greg T. Sueoka**  
**Fenwick & West LLP**  
**Two Palo Alto Square**  
**Palo Alto, CA 94306**  
**U.S.A.**

Telephone	<b>(650) 858-7194</b>	Fax	<b>(650) 494-1417</b>
-----------	-----------------------	-----	-----------------------

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

<b>Name of Sole or First Inventor:</b>		<input type="checkbox"/> A petition has been filed for this unsigned inventor					
Given Name	<b>Dong-Ying</b>	Middle Initial		Family Name	<b>Kuo</b>	Suffix e.g. Jr	
Inventor's Signature					Date		
Residence: City	<b>Pleasanton</b>	State	<b>CA</b>	Country	<b>U.S.A.</b>	Citizenship	<b>U.S.A.</b>
Mailing Address	<b>7857 Olive Court</b>						
Mailing Address							
City	<b>Pleasanton</b>	State	<b>CA</b>	Zip	<b>94588</b>	Country	<b>U.S.A.</b>

☒ Additional inventors are being named on supplemental sheet(s) attached hereto



DECLARATION					ADDITIONAL INVENTOR(S) Supplemental Sheet				
<b>Name of Additional Joint Inventor, if any:</b>					<input type="checkbox"/> A petition has been filed for this unsigned inventor				
Given Name	Zhou			Middle Initial		Family Name	Hong		Suffix e.g. Jr.
Inventor's Signature						Date			
Residence: City	San Jose			State	CA	Country	U.S.A.		Citizenship P.R. China
Mailing Address	311 Checkers Drive, #208								
Mailing Address									
City	San Jose			State	CA	Zip	95133		Country U.S.A.

<b>Name of Additional Joint Inventor, if any:</b>					<input type="checkbox"/> A petition has been filed for this unsigned inventor				
Given Name	Randy			Middle Initial		Family Name	Zhao		Suffix e.g. Jr.
Inventor's Signature						Date			
Residence: City	Fremont			State	CA	Country	94539		Citizenship China
Mailing Address	40338 San Sebastian Place								
Mailing Address									
City	Fremont			State	CA	Zip	94539		Country U.S.A.

<b>Name of Additional Joint Inventor, if any:</b>					<input type="checkbox"/> A petition has been filed for this unsigned inventor				
Given Name	Roger			Middle Initial		Family Name	Niu		Suffix e.g. Jr.
Inventor's Signature						Date			
Residence: City	San Jose			State	CA	Country	U.S.A.		Citizenship U.S.A.
Mailing Address	1061 Harlan Drive								
Mailing Address									
City	San Jose			State	CA	Zip	95129		Country U.S.A.

<b>Name of Additional Joint Inventor, if any:</b>					<input type="checkbox"/> A petition has been filed for this unsigned inventor				
Given Name	Poornachandra			Middle Initial		Family Name	Rao		Suffix e.g. Jr.
Inventor's Signature						Date			
Residence: City	Santa Clara			State	CA	Country	U.S.A.		Citizenship India
Mailing Address	3131 Homestead Road, #9A								
Mailing Address									
City	Santa Clara			State	CA	Zip	95051		Country U.S.A.

☒ Additional inventors are being named on supplemental sheet(s) attached hereto





DECLARATION					ADDITIONAL INVENTOR(S) Supplemental Sheet				
Name of Additional Joint Inventor, if any:					<input type="checkbox"/> A petition has been filed for this unsigned inventor				
Given Name	Lin	Middle Initial		Family Name	Chen	Suffix e.g. Jr.			
Inventor's Signature					Date				
Residence: City	Sunnyvale	State	CA	Country	U.S.A.	Citizenship	P.R. China		
Mailing Address	639 Caliente Drive, #6								
Mailing Address									
City	Sunnyvale	State	CA	Zip	94086	Country	U.S.A.		

Name of Additional Joint Inventor, if any:					<input type="checkbox"/> A petition has been filed for this unsigned inventor				
Given Name	Jeremy	Middle Initial		Family Name	Alves	Suffix e.g. Jr.			
Inventor's Signature					Date				
Residence: City	Mountain View	State	CA	Country	94041	Citizenship	U.S.A.		
Mailing Address	600 Rainbow Drive, #160								
Mailing Address									
City	Mountain View	State	CA	Zip	94041	Country	U.S.A.		

Name of Additional Joint Inventor, if any:					<input type="checkbox"/> A petition has been filed for this unsigned inventor				
Given Name		Middle Initial		Family Name		Suffix e.g. Jr.			
Inventor's Signature					Date				
Residence: City		State		Country		Citizenship			
Mailing Address									
Mailing Address									
City		State		Zip		Country			

Name of Additional Joint Inventor, if any:					<input type="checkbox"/> A petition has been filed for this unsigned inventor				
Given Name		Middle Initial		Family Name		Suffix e.g. Jr.			
Inventor's Signature					Date				
Residence: City		State		Country		Citizenship			
Mailing Address									
Mailing Address									
City		State		Zip		Country			
<input type="checkbox"/> Additional inventors are being named on supplemental sheet(s) attached hereto									

18235/03157/1414

